

Copy from Prior Application

STORAGE ROUTER AND METHOD FOR
PROVIDING VIRTUAL LOCAL STORAGE

TECHNICAL FIELD OF THE INVENTION

This invention relates in general to network storage devices, and more particularly to a storage router and method for providing virtual local storage on remote SCSI storage devices to Fibre Channel devices.

5

BACKGROUND OF THE INVENTION

Typical storage transport mediums provide for a relatively small number of devices to be attached over relatively short distances. One such transport medium is 5 a Small Computer System Interface (SCSI) protocol, the structure and operation of which is generally well known as is described, for example, in the SCSI-1, SCSI-2 and SCSI-3 specifications. High speed serial interconnects provide enhanced capability to attach a large number of 10 high speed devices to a common storage transport medium over large distances. One such serial interconnect is Fibre Channel, the structure and operation of which is described, for example, in *Fibre Channel Physical and 15 Signaling Interface (FC-PH)*, ANSI X3.230 *Fibre Channel Arbitrated Loop (FC-AL)*, and ANSI X3.272 *Fibre Channel Private Loop Direct Attach (FC-PLDA)*.

Conventional computing devices, such as computer workstations, generally access storage locally or through 20 network interconnects. Local storage typically consists of a disk drive, tape drive, CD-ROM drive or other storage device contained within, or locally connected to the workstation. The workstation provides a file system structure, that includes security controls, with access 25 to the local storage device through native low level, block protocols. These protocols map directly to the mechanisms used by the storage device and consist of data requests without security controls. Network interconnects typically provide access for a large number of computing

devices to data storage on a remote network server. The remote network server provides file system structure, access control, and other miscellaneous capabilities that include the network interface. Access to data through 5 the network server is through network protocols that the server must translate into low level requests to the storage device. A workstation with access to the server storage must translate its file system protocols into network protocols that are used to communicate with the 10 server. Consequently, from the perspective of a workstation, or other computing device, seeking to access such server data, the access is much slower than access to data on a local storage device.

SUMMARY OF THE INVENTION

In accordance with the present invention, a storage router and method for providing virtual local storage on remote SCSI storage devices to Fibre Channel devices are disclosed that provide advantages over conventional network storage devices and methods.

According to one aspect of the present invention, a storage router and storage network provide virtual local storage on remote SCSI storage devices to Fibre Channel devices. A plurality of Fibre Channel devices, such as workstations, are connected to a Fibre Channel transport medium, and a plurality of SCSI storage devices are connected to a SCSI bus transport medium. The storage router interfaces between the Fibre Channel transport medium and the SCSI bus transport medium. The storage router maps between the workstations and the SCSI storage devices and implements access controls for storage space on the SCSI storage devices. The storage router then allows access from the workstations to the SCSI storage devices using native low level, block protocol in accordance with the mapping and the access controls.

According to another aspect of the present invention, virtual local storage on remote SCSI storage devices is provided to Fibre Channel devices. A Fibre Channel transport medium and a SCSI bus transport medium are interfaced with. A configuration is maintained for SCSI storage devices connected to the SCSI bus transport medium. The configuration maps between Fibre Channel

devices and the SCSI storage devices and implements access controls for storage space on the SCSI storage devices. Access is then allowed from Fibre Channel initiator devices to SCSI storage devices using native 5 low level, block protocol in accordance with the configuration.

A technical advantage of the present invention is the ability to centralize local storage for networked workstations without any cost of speed or overhead. Each 10 workstation access its virtual local storage as if it work locally connected. Further, the centralized storage devices can be located in a significantly remote position even in excess of ten kilometers as defined by Fibre Channel standards.

Another technical advantage of the present invention 15 is the ability to centrally control and administer storage space for connected users without limiting the speed with which the users can access local data. In addition, global access to data, backups, virus scanning 20 and redundancy can be more easily accomplished by centrally located storage devices.

A further technical advantage of the present invention is providing support for SCSI storage devices as local storage for Fibre Channel hosts. In addition, 25 the present invention helps to provide extended capabilities for Fibre Channel and for management of storage subsystems.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the present invention and the advantages thereof may be acquired by referring to the following description taken in conjunction with the accompanying drawings, in which like reference numbers indicate like features, and wherein:

5

FIGURE 1 is a block diagram of a conventional network that provides storage through a network server;

10 FIGURE 2 is a block diagram of one embodiment of a storage network with a storage router that provides global access and routing;

FIGURE 3 is a block diagram of one embodiment of a storage network with a storage router that provides virtual local storage;

15 FIGURE 4 is a block diagram of one embodiment of the storage router of FIGURE 3; and

FIGURE 5 is a block diagram of one embodiment of data flow within the storage router of FIGURE 4.

DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 is a block diagram of a conventional network, indicated generally at 10, that provides access to storage through a network server. As shown, network 10 includes a plurality of workstations 12 interconnected with a network server 14 via a network transport medium 16. Each workstation 12 can generally comprise a processor, memory, input/output devices, storage devices and a network adapter as well as other common computer components. Network server 14 uses a SCSI bus 18 as a storage transport medium to interconnect with a plurality of storage devices 20 (tape drives, disk drives, etc.). In the embodiment of FIGURE 1, network transport medium 16 is an network connection and storage devices 20 comprise hard disk drives, although there are numerous alternate transport mediums and storage devices.

In network 10, each workstation 12 has access to its local storage device as well as network access to data on storage devices 20. The access to a local storage device is typically through native low level, block protocols. On the other hand, access by a workstation 12 to storage devices 20 requires the participation of network server 14 which implements a file system and transfers data to workstations 12 only through high level file system protocols. Only network server 14 communicates with storage devices 20 via native low level, block protocols. Consequently, the network access by workstations 12 through network server 14 is slow with respect to their

access to local storage. In network 10, it can also be a logistical problem to centrally manage and administer local data distributed across an organization, including accomplishing tasks such as backups, virus scanning and 5 redundancy.

FIGURE 2 is a block diagram of one embodiment of a storage network, indicated generally at 30, with a storage router that provides global access and routing. This environment is significantly different from that of 10 FIGURE 1 in that there is no network server involved. In FIGURE 2, a Fibre Channel high speed serial transport 32 interconnects a plurality of workstations 36 and storage devices 38. A SCSI bus storage transport medium interconnects workstations 40 and storage devices 42. A 15 storage router 44 then serves to interconnect these mediums and provide devices on either medium global, transparent access to devices on the other medium. Storage router 44 routes requests from initiator devices on one medium to target devices on the other medium and 20 routes data between the target and the initiator. Storage router 44 can allow initiators and targets to be on either side. In this manner, storage router 44 enhances the functionality of Fibre Channel 32 by providing access, for example, to legacy SCSI storage 25 devices on SCSI bus 34. In the embodiment of FIGURE 2, the operation of storage router 44 can be managed by a management station 46 connected to the storage router via a direct serial connection.

In storage network 30, any workstation 36 or workstation 40 can access any storage device 38 or storage device 42 through native low level, block protocols, and vice versa. This functionality is enabled by storage router 44 which routes requests and data as a generic transport between Fibre Channel 32 and SCSI bus 34. Storage router 44 uses tables to map devices from one medium to the other and distributes requests and data across Fibre Channel 32 and SCSI bus 34 without any security access controls. Although this extension of the high speed serial interconnect provided by Fibre Channel 32 is beneficial, it is desirable to provide security controls in addition to extended access to storage devices through a native low level, block protocol.

FIGURE 3 is a block diagram of one embodiment of a storage network, indicated generally at 50, with a storage router that provides virtual local storage. Similar to that of FIGURE 2, storage network 50 includes a Fibre Channel high speed serial interconnect 52 and a SCSI bus 54 bridged by a storage router 56. Storage router 56 of FIGURE 3 provides for a large number of workstations 58 to be interconnected on a common storage transport and to access common storage devices 60, 62 and 64 through native low level, block protocols.

According to the present invention, storage router 56 has enhanced functionality to implement security controls and routing such that each workstation 58 can have access to a specific subset of the overall data

5 stored in storage devices 60, 62 and 64. This specific subset of data has the appearance and characteristics of local storage and is referred to herein as virtual local storage. Storage router 56 allows the configuration and
10 modification of the storage allocated to each attached workstation 58 through the use of mapping tables or other mapping techniques.

15 As shown in FIGURE 3, for example, storage device 60 can be configured to provide global data 65 which can be accessed by all workstations 58. Storage device 62 can be configured to provide partitioned subsets 66, 68, 70 and 72, where each partition is allocated to one of the workstations 58 (workstations A, B, C and D). These subsets 66, 68, 70 and 72 can only be accessed by the associated workstation 58 and appear to the associated workstation 58 as local storage accessed using native low level, block protocols. Similarly, storage device 64 can be allocated as storage for the remaining workstation 58 (workstation E).

20 Storage router 56 combines access control with routing such that each workstation 58 has controlled access to only the specified partition of storage device 62 which forms virtual local storage for the workstation 58. This access control allows security control for the specified data partitions. Storage router 56 allows this allocation of storage devices 60, 62 and 64 to be managed by a management station 76. Management station 76 can connect directly to storage router 56 via a direct

connection or, alternately, can interface with storage router 56 through either Fibre Channel 52 or SCSI bus 54. In the latter case, management station 76 can be a workstation or other computing device with special rights 5 such that storage router 56 allows access to mapping tables and shows storage devices 60, 62 and 64 as they exist physically rather than as they have been allocated.

The environment of FIGURE 3 extends the concept of a single workstation having locally connected storage 10 devices to a storage network 50 in which workstations 58 are provided virtual local storage in a manner transparent to workstations 58. Storage router 56 provides centralized control of what each workstation 58 sees as its local drive, as well as what data it sees as 15 global data accessible by other workstations 58.

Consequently, the storage space considered by the workstation 58 to be its local storage is actually a partition (i.e., logical storage definition) of a physically remote storage device 60, 62 or 64 connected 20 through storage router 56. This means that similar requests from workstations 58 for access to their local storage devices produce different accesses to the storage space on storage devices 60, 62 and 64. Further, no access from a workstation 58 is allowed to the virtual 25 local storage of another workstation 58.

The collective storage provided by storage devices 60, 62 and 64 can have blocks allocated by programming means within storage router 56. To accomplish this

function, storage router 56 can include routing tables and security controls that define storage allocation for each workstation 58. The advantages provided by implementing virtual local storage in centralized storage 5 devices include the ability to do collective backups and other collective administrative functions more easily. This is accomplished without limiting the performance of workstations 58 because storage access involves native 10 low level, block protocols and does not involve the overhead of high level protocols and file systems required by network servers.

FIGURE 4 is a block diagram of one embodiment of storage router 56 of FIGURE 3. Storage router 56 can comprise a Fibre Channel controller 80 that interfaces 15 with Fibre Channel 52 and a SCSI controller 82 that interfaces with SCSI bus 54. A buffer 84 provides memory work space and is connected to both Fibre Channel controller 80 and to SCSI controller 82. A supervisor unit 86 is connected to Fibre Channel controller 80, SCSI 20 controller 82 and buffer 84. Supervisor unit 86 comprises a microprocessor for controlling operation of storage router 56 and to handle mapping and security access for requests between Fibre Channel 52 and SCSI bus 54.

FIGURE 5 is a block diagram of one embodiment of 25 data flow within storage router 56 of FIGURE 4. As shown, data from Fibre Channel 52 is processed by a Fibre Channel (FC) protocol unit 88 and placed in a FIFO queue

90. A direct memory access (DMA) interface 92 then takes data out of FIFO queue 90 and places it in buffer 84. Supervisor unit 86 processes the data in buffer 84 as represented by supervisor processing 93. This processing 5 involves mapping between Fibre Channel 52 and SCSI bus 54 and applying access controls and routing functions. A DMA interface 94 then pulls data from buffer 84 and places it into a buffer 96. A SCSI protocol unit 98 pulls data from buffer 96 and communicates the data on 10 SCSI bus 54. Data flow in the reverse direction, from SCSI bus 54 to Fibre Channel 52, is accomplished in a reverse manner.

The storage router of the present invention is a bridge device that connects a Fibre Channel link directly 15 to a SCSI bus and enables the exchange of SCSI command set information between application clients on SCSI bus devices and the Fibre Channel links. Further, the storage router applies access controls such that virtual local storage can be established in remote SCSI storage 20 devices for workstations on the Fibre Channel link. In one embodiment, the storage router provides a connection for Fibre Channel links running the SCSI Fibre Channel Protocol (FCP) to legacy SCSI devices attached to a SCSI bus. The Fibre Channel topology is typically an 25 Arbitrated Loop (FC_AL).

In part, the storage router enables a migration path to Fibre Channel based, serial SCSI networks by providing connectivity for legacy SCSI bus devices. The storage

router can be attached to a Fibre Channel Arbitrated Loop and a SCSI bus to support a number of SCSI devices.

Using configuration settings, the storage router can make the SCSI bus devices available on the Fibre Channel

5 network as FCP logical units. Once the configuration is defined, operation of the storage router is transparent to application clients. In this manner, the storage router can form an integral part of the migration to new Fibre Channel based networks while providing a means to
10 continue using legacy SCSI devices.

In one implementation (not shown), the storage router can be a rack mount or free standing device with an internal power supply. The storage router can have a Fibre Channel and SCSI port, and a standard, detachable power cord can be used, the FC connector can be a copper DB9 connector, and the SCSI connector can be a 68-pin type. Additional modular jacks can be provided for a serial port and a 802.3 10BaseT port, i.e. twisted pair Ethernet, for management access. The SCSI port of the
15 storage router can support SCSI direct and sequential access target devices and can support SCSI initiators, as well. The Fibre Channel port can interface to SCSI-3 FCP enabled devices and initiators.

20 To accomplish its functionality, one implementation of the storage router uses: a Fibre Channel interface based on the HEWLETT-PACKARD TACHYON HPFC-5000 controller and a GLM media interface; an Intel 80960RP processor, incorporating independent data and program memory spaces,

and associated logic required to implement a stand alone processing system; and a serial port for debug and system configuration. Further, this implementation includes a SCSI interface supporting Fast-20 based on the SYMBIOS 5 53C8xx series SCSI controllers, and an operating system based upon the WIND RIVERS SYSTEMS VXWORKS or IXWORKS kernel, as determined by design. In addition, the storage router includes software as required to control basic functions of the various elements, and to provide 10 appropriate translations between the FC and SCSI protocols.

The storage router has various modes of operation that are possible between FC and SCSI target and initiator combinations. These modes are: FC Initiator to 15 SCSI Target; SCSI Initiator to FC Target; SCSI Initiator to SCSI Target; and FC Initiator to FC Target. The first two modes can be supported concurrently in a single storage router device are discussed briefly below. The third mode can involve two storage router devices 20 back to back and can serve primarily as a device to extend the physical distance beyond that possible via a direct SCSI connection. The last mode can be used to carry FC protocols encapsulated on other transmission technologies (e.g. ATM, SONET), or to act as a bridge 25 between two FC loops (e.g. as a two port fabric).

The FC Initiator to SCSI Target mode provides for the basic configuration of a server using Fibre Channel to communicate with SCSI targets. This mode requires

that a host system have an FC attached device and associated device drivers and software to generate SCSI-3 FCP requests. This system acts as an initiator using the storage router to communicate with SCSI target devices. 5 The SCSI devices supported can include SCSI-2 compliant direct or sequential access (disk or tape) devices. The storage router serves to translate command and status information and transfer data between SCSI-3 FCP and SCSI-2, allowing the use of standard SCSI-2 devices in a 10 Fibre Channel environment.

The SCSI Initiator to FC Target mode provides for the configuration of a server using SCSI-2 to communicate with Fibre Channel targets. This mode requires that a host system has a SCSI-2 interface and driver software to 15 control SCSI-2 target devices. The storage router will connect to the SCSI-2 bus and respond as a target to multiple target IDs. Configuration information is required to identify the target IDs to which the bridge will respond on the SCSI-2 bus. The storage router then 20 translates the SCSI-2 requests to SCSI-3 FCP requests, allowing the use of FC devices with a SCSI host system. This will also allow features such as a tape device acting as an initiator on the SCSI bus to provide full support for this type of SCSI device.

25 In general, user configuration of the storage router will be needed to support various functional modes of operation. Configuration can be modified, for example, through a serial port or through an Ethernet port via

SNMP (simple network management protocol) or a Telnet session. Specifically, SNMP manageability can be provided via an 802.3 Ethernet interface. This can provide for configuration changes as well as providing 5 statistics and error information. Configuration can also be performed via TELNET or RS-232 interfaces with menu driven command interfaces. Configuration information can be stored in a segment of flash memory and can be retained across resets and power off cycles. Password 10 protection can also be provided.

In the first two modes of operation, addressing information is needed to map from FC addressing to SCSI addressing and vice versa. This can be 'hard' configuration data, due to the need for address 15 information to be maintained across initialization and partial reconfigurations of the Fibre Channel address space. In an arbitrated loop configuration, user configured addresses will be needed for AL_PAs in order to insure that known addresses are provided between loop 20 reconfigurations.

With respect to addressing, FCP and SCSI 2 systems employ different methods of addressing target devices. Additionally, the inclusion of a storage router means that a method of translating device IDs needs to be 25 implemented. In addition, the storage router can respond to commands without passing the commands through to the opposite interface. This can be implemented to allow all generic FCP and SCSI commands to pass through the storage

router to address attached devices, but allow for configuration and diagnostics to be performed directly on the storage router through the FC and SCSI interfaces.

Management commands are those intended to be processed by the storage router controller directly. This may include diagnostic, mode, and log commands as well as other vendor-specific commands. These commands can be received and processed by both the FCP and SCSI interfaces, but are not typically bridged to the opposite interface. These commands may also have side effects on the operation of the storage router, and cause other storage router operations to change or terminate.

A primary method of addressing management commands through the FCP and SCSI interfaces can be through peripheral device type addressing. For example, the storage router can respond to all operations addressed to logical unit (LUN) zero as a controller device. Commands that the storage router will support can include INQUIRY as well as vendor-specific management commands. These are to be generally consistent with SCC standard commands.

The SCSI bus is capable of establishing bus connections between targets. These targets may internally address logical units. Thus, the prioritized addressing scheme used by SCSI subsystems can be represented as follows: BUS:TARGET:LOGICAL UNIT. The BUS identification is intrinsic in the configuration, as a SCSI initiator is attached to only one bus. Target

addressing is handled by bus arbitration from information provided to the arbitrating device. Target addresses are assigned to SCSI devices directly, though some means of configuration, such as a hardware jumper, switch setting, 5 or device specific software configuration. As such, the SCSI protocol provides only logical unit addressing within the Identify message. Bus and target information is implied by the established connection.

10 Fibre Channel devices within a fabric are addressed by a unique port identifier. This identifier is assigned to a port during certain well-defined states of the FC protocol. Individual ports are allowed to arbitrate for a known, user defined address. If such an address is not provided, or if arbitration for a particular user address 15 fails, the port is assigned a unique address by the FC protocol. This address is generally not guaranteed to be unique between instances. Various scenarios exist where the AL_PA of a device will change, either after power cycle or loop reconfiguration.

20 The FC protocol also provides a logical unit address field within command structures to provide addressing to devices internal to a port. The FCP_CMD payload specifies an eight byte LUN field. Subsequent identification of the exchange between devices is 25 provided by the FQXID (Fully Qualified Exchange ID).

FC ports can be required to have specific addresses assigned. Although basic functionality is not dependent on this, changes in the loop configuration could result

in disk targets changing identifiers with the potential risk of data corruption or loss. This configuration can be straightforward, and can consist of providing the device a loop-unique ID (AL_PA) in the range of "01h" to 5 "EFh." Storage routers could be shipped with a default value with the assumption that most configurations will be using single storage routers and no other devices requesting the present ID. This would provide a minimum amount of initial configuration to the system
10 administrator. Alternately, storage routers could be defaulted to assume any address so that configurations requiring multiple storage routers on a loop would not require that the administrator assign a unique ID to the additional storage routers.

15 Address translation is needed where commands are issued in the cases FC Initiator to SCSI Target and SCSI Initiator to FC Target. Target responses are qualified by the FQXID and will retain the translation acquired at the beginning of the exchange. This prevents
20 configuration changes occurring during the course of execution of a command from causing data or state information to be inadvertently misdirected. Configuration can be required in cases of SCSI Initiator to FC Target, as discovery may not effectively allow for
25 FCP targets to consistently be found. This is due to an FC arbitrated loop supporting addressing of a larger number of devices than a SCSI bus and the possibility of FC devices changing their AL-PA due to device insertion

or other loop initialization.

In the direct method, the translation to BUS:TARGET:LUN of the SCSI address information will be direct. That is, the values represented in the FCP LUN field will directly map to the values in effect on the SCSI bus. This provides a clean translation and does not require SCSI bus discovery. It also allows devices to be dynamically added to the SCSI bus without modifying the address map. It may not allow for complete discovery by FCP initiator devices, as gaps between device addresses may halt the discovery process. Legacy SCSI device drivers typically halt discovery on a target device at the first unoccupied LUN, and proceed to the next target. This would lead to some devices not being discovered. However, this allows for hot plugged devices and other changes to the loop addressing.

In the ordered method, ordered translation requires that the storage router perform discovery on reset, and collapses the addresses on the SCSI bus to sequential FCP LUN values. Thus, the FCP LUN values 0-N can represent N+1 SCSI devices, regardless of SCSI address values, in the order in which they are isolated during the SCSI discovery process. This would allow the FCP initiator discovery process to identify all mapped SCSI devices without further configuration. This has the limitation that hot-plugged devices will not be identified until the next reset cycle. In this case, the address may also be altered as well.

In addition to addressing, according to the present invention, the storage router provides configuration and access controls that cause certain requests from FC Initiators to be directed to assigned virtual local storage partitioned on SCSI storage devices. For example, the same request for LUN 0 (local storage) by two different FC Initiators can be directed to two separate subsets of storage. The storage router can use tables to map, for each initiator, what storage access is available and what partition is being addressed by a particular request. In this manner, the storage space provided by SCSI storage devices can be allocated to FC initiators to provide virtual local storage as well as to create any other desired configuration for secured access.

Although the present invention has been described in detail, it should be understood that various changes, substitutions, and alterations can be made hereto without departing from the spirit and scope of the invention as defined by the appended claims.